



DP-420^{Q&As}

Designing and Implementing Cloud-Native Applications Using Microsoft Azure Cosmos DB

Pass Microsoft DP-420 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/dp-420.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Cosmos DB Core (SQL) API account named account 1 that uses autoscale throughput.

You need to run an Azure function when the normalized request units per second for a container in account1 exceeds a specific value.

Solution: You configure an Azure Monitor alert to trigger the function.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

You can set up alerts from the Azure Cosmos DB pane or the Azure Monitor service in the Azure portal.

Note: Alerts are used to set up recurring tests to monitor the availability and responsiveness of your Azure Cosmos DB resources. Alerts can send you a notification in the form of an email, or execute an Azure Function when one of your metrics reaches the threshold or if a specific event is logged in the activity log.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/create-alerts>

QUESTION 2

HOTSPOT

You are creating a database in an Azure Cosmos DB Core (SQL) API account. The database will be used by an application that will provide users with the ability to share online posts. Users will also be able to submit comments on other

users\' posts.

You need to store the data shown in the following table.

Type	Description
Users	Information about a user who will use the application
Posts	Text of up to 1,000 characters that a user will share with other users
Comments	Text of up to 280 characters that users will submit as a comment on a post
Interests	Information about a user\'s interests

The application has the following characteristics:



1.

Users can submit an unlimited number of posts.

2.

The average number of posts submitted by a user will be more than 1,000.

3.

Posts can have an unlimited number of comments from different users.

4.

The average number of comments per post will be 100, but many posts will exceed 1,000 comments.

5.

Users will be limited to having a maximum of 20 interests.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
If you embed the posts data into the users data instead of creating a separate document for each post, you will increase the write operation costs for new posts	<input type="radio"/>	<input type="radio"/>
If you embed the comments data into the posts data instead of creating a separate document for each comment you will increase the write operation costs for new comments	<input type="radio"/>	<input type="radio"/>
If you embed the interests data into the users data instead of creating a separate document for each interest, you will increase the read operation costs for displaying the users and their associated interests	<input type="radio"/>	<input type="radio"/>

Correct Answer:



Answer Area

Statements

Yes

No

If you embed the posts data into the users data instead of creating a separate document for each post, you will increase the write operation costs for new posts

☒☐

If you embed the comments data into the posts data instead of creating a separate document for each comment you will increase the write operation costs for new comments

☒☐

If you embed the interests data into the users data instead of creating a separate document for each interest, you will increase the read operation costs for displaying the users and their associated interests

☐☒

Box 1: Yes

Non-relational data increases write costs, but can decrease read costs.

Box 2: Yes

Non-relational data increases write costs, but can decrease read costs.

Box 3: No

Non-relational data increases write costs, but can decrease read costs.

QUESTION 3

The settings for a container in an Azure Cosmos DB Core (SQL) API account are configured as shown in the following exhibit.



Settings

Indexing Policy

Time to Live

- ☐ Off
- ☒ On (no default)
- ☐ On

Geospatial Configuration

- ☒ Geography
- ☐ Geometry

Partition key

Which statement describes the configuration of the container?

- A. All items will be deleted after one year.
- B. Items stored in the collection will be retained always, regardless of the items time to live value.
- C. Items stored in the collection will expire only if the item has a time to live value.
- D. All items will be deleted after one hour.

Correct Answer: C

When DefaultTimeToLive is -1 then your Time to Live setting is On (No default)

Time to Live on a container, if present and the value is set to "-1", it is equal to infinity, and items don't expire by default.

Time to Live on an item:

This Property is applicable only if DefaultTimeToLive is present and it is not set to null for the parent container.

If present, it overrides the DefaultTimeToLive value of the parent container.

Reference:



<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/time-to-live>

QUESTION 4

You have an Azure Cosmos DB for NoSQL account configured for global distribution across four regions.

At connection time, the SQL SDK needs to identify the optimal endpoint for reading and writing.

Which two factors can influence the SDK? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. the consistency level in RequestOptions
- B. regional latency
- C. the default consistency level
- D. the PreferredLocations configuration
- E. a region being available

Correct Answer: DE

Explanation:

Connecting to a preferred region using the API for NoSQL

The SDKs accept an optional parameter PreferredLocations that is an ordered list of Azure regions.

The SDK will automatically send all writes to the current write region. All reads will be sent to the first available region in the preferred locations list. If the request fails, the client will fail down the list to the next region.

The SDK will only attempt to read from the regions specified in preferred locations. So, for example, if the Azure Cosmos DB account is available in four regions, but the client only specifies two read(non-write) regions within the

PreferredLocations, then no reads will be served out of the read region that is not specified in PreferredLocations. If the read regions specified in the PreferredLocations list are not available, reads will be served out of write region.

Reference: <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/tutorial-global-distribution>

QUESTION 5

You have operational data in an Azure Cosmos DB for NoSQL database.

Database users report that the performance of the database degrades significantly when a business analytics team runs large Apache Spark-based queries against the database.

You need to reduce the impact that running the Spark-based queries has on the database users.

What should you implement?

- A. Azure Synapse Link



B. a default consistency level of Consistent Prefix

C. a default consistency level of Strong

D. the Spark connector

Correct Answer: A

[DP-420 PDF Dumps](#)

[DP-420 VCE Dumps](#)

[DP-420 Study Guide](#)